

Face frontalization for Alignment and Recognition

Christos Sagonas*, Yannis Panagakis*, Stefanos Zafeiriou*, and Maja Pantic*[†]

*Department of Computing,

Imperial College London,

180 Queens Gate,

London SW7 2AZ, U.K.

[†]EEMCS,

University of Twente,

Drienerlolaan 5,

7522 NB Enschede, The Netherlands

{c.sagonas, i.panagakis, s.zafeiriou, m.pantic}@imperial.ac.uk

Abstract—Recently, it was shown that excellent results can be achieved in both face landmark localization and pose-invariant face recognition. These breakthroughs are attributed to the efforts of the community to manually annotate facial images in many different poses and to collect 3D faces data. In this paper, we propose a novel method for joint face landmark localization and frontal face reconstruction (pose correction) using a small set of frontal images only. By observing that the frontal facial image is the one with the minimum rank from all different poses we formulate an appropriate model which is able to jointly recover the facial landmarks as well as the frontalized version of the face. To this end, a suitable optimization problem, involving the minimization of the nuclear norm and the matrix ℓ_1 norm, is solved. The proposed method is assessed in frontal face reconstruction (pose correction), face landmark localization, and pose-invariant face recognition and verification by conducting experiments on 6 facial images databases. The experimental results demonstrate the effectiveness of the proposed method.

I. INTRODUCTION

Face analysis is one of the most popular computer vision problems. Important topics in face analysis include generic face alignment [5], [39] and automatic face recognition [4], [40]. These problems have been considered as separate, both creating a wealth of scientific research in Computer Vision. In particular state-of-the-art face alignment and landmark localization methods [5], [39] model the problem discriminatively by capitalizing on the availability of annotated, with regards to facial landmarks, data [30], [31]. Unfortunately, the annotation of facial landmark is laborious, expensive, and time consuming process. This is more evident in cases where the face is not in frontal pose and some facial features and the boundary are neither visible nor well-defined.¹

Existing methods for face alignment can be roughly divided into two main categories: (a) *Holistic* methods which use the whole texture of face as representation, and (b) *part-based* methods which represent the face by using a set of local image patches extracted around of the predefined landmark points. The most-well known methods from the first category are the Active Appearance Models (AAMs) [12], [25], [35] and the 3D Deformable Models (3DMs) [8]. In the second category, methods such as the Active Shape Models (ASMs) [13] and the Constrained Local Models (CLMs) [14], [32] are included. Many of the above mentioned face alignment methods have achieved state-of-the-art results (e.g., [5], [39]) in facial landmark localization under in-the-wild conditions

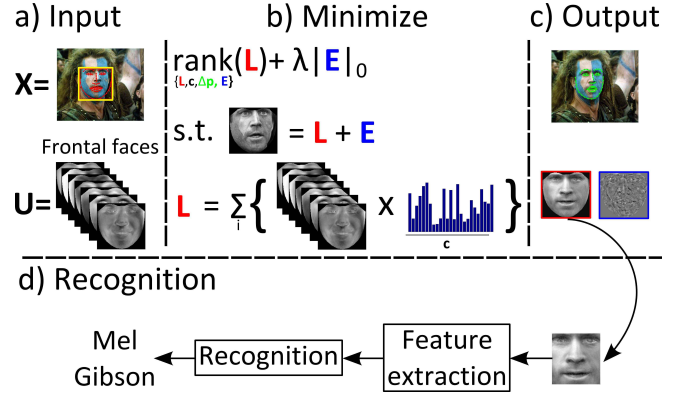


Fig. 1. Flowchart of the proposed method: a) Given an input image, the results from a detector, and a statistical model U , built on frontal images only, b) a constrained low-rank minimization problem is solved. c) Face alignment and frontal view reconstruction are performed simultaneously. Finally, d) face recognition is succeeded using the frontalized image.

but they are trained on many annotated samples from various poses.

On the other hand, in the majority of face recognition systems the first and arguably, most important, step of face alignment is taken for granted using off the self methods [38], [39]. Even in the recent state-of-the-art face recognition methods, where millions of image are used to train feature extractors and classifiers, the pivotal step that increases their performance is that of face alignment [34], [40]. In such cases, the alignment step is very elaborate requiring to both locate landmarks and use 3D face models for pose corrections. In general, 3D model-based methods have high recognition accuracy due to the incorporation of the 3D model. However, such methods cannot be widely applied since they require: (a) a method for accurate landmark localization in various poses, (b) to fit learned 3D model of faces, which is expensive to build, and (c) to develop robust image warping algorithms to reconstruct the frontal image [34]. A recent approach that does not require a 3D model but only a small set of landmarks is presented in [18]. This method aims to reconstruct the virtual view of a non-frontal image by employing Markov Random Field. The main drawback of the aforementioned method is that for each non-frontal image an exhaustively batch-based alignment algorithm trained on frontal patches is applied. Clearly, such a procedure is time consuming.

In this paper, motivated by the observation that the rank of a frontal facial image, due to the approximately structure

¹From experience we know that annotation of facial image with poses take in many cases twice the time compared with frontal poses.




Fig. 2. The average value of nuclear norm computed based on neutral images of twenty subjects from MultiPIE database under poses $-30^0 : 30^0$. The initial images and the warped ones of a subject are also depicted.

of human face, is much smaller than the rank of facial images in other poses, we propose a unified framework for joint face frontalization (pose correction), landmark localization and (single sample) pose invariant face recognition. We show this can be achieved by using a model built from **frontal images only**. To validate the above observation ‘Neutral’ images of twenty objects from MultiPIE database under poses $-30^0 : 30^0$ were warped into a reference frame and the nuclear norm (convex surrogate of the rank) of each shape-free texture was computed. In Fig. 2 the average value of the nuclear norm for the different poses is depicted. Clearly, the frontal pose has the smallest nuclear norm value compared the corresponding value of the rest of the poses. The flowchart of the proposed method (coined as FAR - *Face frontalization for Alignment and Recognition*) is depicted in Fig 1.

The most closely related work to the proposed method is the Transform Invariant Low-rank Textures (TILT) method [43]. In TILT, texture rectification is obtained by applying a global affine transformation onto a low-rank term, modelling the texture. By blindly imposing low-rank constraints without regularization, for non-rigid alignment opposite effects may occur. We applied the TILT with a non-rigid facial shape model and its performance was very poor as it can be observed in Fig. 3. Recently, it was demonstrated [11], [29], that non-rigid deformable models cannot be straightforward combined with optimization problems [27] that involve low-rank terms without a proper regularization. To overcome this and ensure




Fig. 3. Results produced by the (a) non-rigid TILT and (b) FAR.

that unnatural faces will be not created, a model of frontal

images is employed. In that sense, our method can be seen as a deformable TILT model regularized within a frontal face subspace.

Summarizing, the contributions of the paper are:

- Technical contribution: We develop a joint landmark localization and face frontalization method by proposing a deformable and appropriately regularized TILT.
- Applications in computer vision:
 - 1) To the best of our knowledge this is the first generic landmark localization method which achieves state-of-the-art results using a model of **frontal images only**.
 - 2) It is possible to improve or match the state-of-the-art in unconstrained face recognition using only frontal faces and simple features for classification unlike other complex feature extraction procedures e.g., [33].²

II. NOTATION AND PRELIMINARIES

Throughout the paper, scalars are denoted by lower-case letters, vectors (matrices) are denoted by lower-case (upper-case) boldface letters i.e., \mathbf{x} , (\mathbf{X}). \mathbf{I} denotes the identity matrix of compatible dimensions. The i th column of \mathbf{X} is denoted by \mathbf{x}_i . A vector $\mathbf{x} \in \mathbb{R}^{m \cdot n}$ (matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$) is reshaped into a matrix (vector) via the reshape operator : $\mathcal{R}_{m \times n}(\mathbf{x}) = \mathbf{X} \in \mathbb{R}^{m \times n}$, ($\text{vec}(\mathbf{X}) = \mathbf{x} \in \mathbb{R}^{m \cdot n}$).

The ℓ_1 and the ℓ_2 norms of \mathbf{x} are defined as $\|\mathbf{x}\|_1 = \sum_i |x_i|$ and $\|\mathbf{x}\|_2 = \sqrt{\sum_i x_i^2}$, respectively. The matrix ℓ_1 norm is defined as $\|\mathbf{X}\|_1 = \sum_i \sum_j |x_{ij}|$, where $|\cdot|$ denotes the absolute value operator. The Frobenius norm is defined as $\|\mathbf{X}\|_F = \sqrt{\sum_i \sum_j x_{ij}^2}$, and the nuclear norm of \mathbf{X} (i.e., the sum of singular values of a matrix) is denoted by $\|\mathbf{X}\|_*$. \mathbf{X}^T is the transpose of \mathbf{X} . If \mathbf{X} is a square matrix, \mathbf{X}^{-1} is its inverse, provided that the inverse matrix exists.

The warp function $\mathbf{x}(\mathcal{W}(\mathbf{p}))$ ($\mathbf{X}(\mathcal{W}(\mathbf{p}))$) denotes the warping of 2D coordinates arranged as vector (matrix) by a warp parameter vector $\mathbf{p} \in \mathbb{R}^v$, where v is the number of the warping parameters, back to reference coordinate system. To simplify the notation $\mathbf{x}(\mathbf{p})$ ($\mathbf{X}(\mathbf{p})$) will be used throughout the paper instead of $\mathbf{x}(\mathcal{W}(\mathbf{p}))$ ($\mathbf{X}(\mathcal{W}(\mathbf{p}))$).

III. PROPOSED METHOD

A. Problem formulation

In this paper, the goal is to recover the frontal view (i.e., $\mathbf{L} \in \mathbb{R}^{m \times n}$) of a warped facial image (i.e., $\mathbf{X}(\mathbf{p}) \in \mathbb{R}^{m \times n}$) which is possibly corrupted by sparse error of large magnitude. Such sparse errors indicate that only a small fraction of the image pixels may be corrupted by non-Gaussian noise and occlusions. In particular, based on the observation that the frontal view of a face lies onto a low-rank subspace (please refer to Fig. 2), it can be expressed as a linear combination of a small number of orthonormal learned basis (i.e. $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k] \in \mathbb{R}^{f \times k}$, $\mathbf{U}^T \mathbf{U} = \mathbf{I}$) that span a generic

²We note that we refer to the restricted protocol of the LFW [20] and not to the unrestricted which unfortunately we cannot compete since we do not have access to millions of annotated faces.

clean frontal face subspace, that is $\mathbf{L} = \sum_{i=1}^k \mathcal{R}_{m \times n}(\mathbf{u}_i) c_i$. Therefore, a warped corrupted image is written as:

$$\mathbf{X}(\mathbf{p}) = \mathbf{L} + \mathbf{E} = \sum_{i=1}^k \mathcal{R}_{m \times n}(\mathbf{u}_i) c_i + \mathbf{E}, \quad (1)$$

where \mathbf{E} is the sparse error matrix.

To match nicely the specifications of the frontal image and the sparse error one can find the low-rank frontal image, the linear combination coefficients, the increments of warp parameters, and the error matrix by solving:

$$\begin{aligned} & \underset{\mathbf{L}, \mathbf{E}, \mathbf{c}, \Delta \mathbf{p}}{\operatorname{argmin}} \quad \|\mathbf{L}\|_* + \lambda \|\mathbf{E}\|_1 \\ \text{s.t.} \quad & \mathbf{X}(\mathbf{p}) = \mathbf{L} + \mathbf{E}, \quad \mathbf{L} = \sum_{i=1}^k \mathcal{R}_{m \times n}(\mathbf{u}_i) c_i, \end{aligned} \quad (2)$$

where the nuclear norm $\|\mathbf{L}\|_*$ and the ℓ_1 norm $\|\mathbf{E}\|_1$ are utilized to promote low-rank on \mathbf{L} and sparsity in \mathbf{E} , while λ is a positive parameter balancing the norms. The nuclear and ℓ_1 norms are the closest convex surrogates to the natural criteria of rank [16] and cardinality [15] which are NP-hard in general to optimize [26], [37]. However, (2) is difficult to be solved due to the non-linearity of the constraint $\mathbf{X}(\mathbf{p}) = \mathbf{L} + \mathbf{E}$. To remedy this, a first order Taylor linear approximation is applied on the vectorized form of the constraint: $\mathbf{x}(\mathbf{p} + \Delta \mathbf{p}) \approx \mathbf{x}(\mathbf{p}) + \mathbf{J}(\mathbf{p}) \Delta \mathbf{p}$, where $\operatorname{vec}(\mathbf{X}(\mathbf{p})) = \operatorname{vec}(\mathbf{L} + \mathbf{E}) = \mathbf{U} \mathbf{c} + \mathbf{e} = \mathbf{x}(\mathbf{p})$ and $\mathbf{J}(\mathbf{p}) = \nabla \mathbf{x}(\mathbf{p}) \frac{\partial \mathbf{W}}{\partial \mathbf{p}}$ is the Jacobian matrix with the steepest descent images as its columns. Consequently, (2) is written as:

$$\begin{aligned} & \underset{\mathbf{L}, \mathbf{E}, \mathbf{c}, \mathbf{p}}{\operatorname{argmin}} \quad \|\mathbf{L}\|_* + \lambda \|\mathbf{E}\|_1 \\ \text{s.t.} \quad & h_1(\Delta \mathbf{p}, \mathbf{c}, \mathbf{e}) = 0, \quad h_2(\mathbf{L}, \mathbf{c}) = 0, \end{aligned} \quad (3)$$

where $h_1(\Delta \mathbf{p}, \mathbf{c}, \mathbf{e}) = \mathbf{x}(\mathbf{p}) + \mathbf{J}(\mathbf{p}) \Delta \mathbf{p} - \mathbf{U} \mathbf{c} - \mathbf{e}$ and $h_2(\mathbf{L}, \mathbf{c}) = \mathbf{L} - \sum_{i=1}^k \mathcal{R}_{m \times n}(\mathbf{u}_i) c_i$.

B. Optimization

To solve (3), the *augmented* Lagrangian [7] is introduced:

$$\begin{aligned} \mathcal{L}(\mathbf{L}, \mathbf{c}, \Delta \mathbf{p}, \mathbf{e}, \mathbf{a}, \mathbf{B}) = & \|\mathbf{L}\|_* + \lambda \|\mathbf{e}\|_1 + \mathbf{a}^T (h_1(\Delta \mathbf{p}, \mathbf{c}, \mathbf{e})) \\ & + \operatorname{tr}(\mathbf{B}^T (h_2(\mathbf{L}, \mathbf{c}))) + \frac{\mu}{2} (\|h_1(\Delta \mathbf{p}, \mathbf{c}, \mathbf{e})\|_2^2 + \|h_2(\mathbf{L}, \mathbf{c})\|_F^2), \end{aligned} \quad (4)$$

where \mathbf{a} and \mathbf{B} the Lagrange multipliers for the equality constraints in (3) and $\mu > 0$ is a penalty parameter. By employing the Alternating Directions Method (ADM) of multipliers [7], (3) is solved by minimizing (4) with respect to each variable in an alternating fashion and finally the Lagrange multipliers are updated at each iteration as outlined in Algorithm 1 its derivation is provided next.

Let t be the iteration index. Given the $\mathbf{L}_t, \mathbf{c}_t, \Delta \mathbf{p}_t, \mathbf{e}_t, \mathbf{a}_t$, and \mathbf{B}_t , the updates are computed by solving the following

sub-problems:

$$\mathbf{L}_{t+1} = \underset{\mathbf{L}_t}{\operatorname{argmin}} \mathcal{L}(\mathbf{L}_t, \mathbf{c}_t, \Delta \mathbf{p}_t, \mathbf{e}_t, \mathbf{a}_t, \mathbf{B}_t) \quad (5)$$

$$\mathbf{c}_{i,t+1} = \underset{\mathbf{c}_{i,t}}{\operatorname{argmin}} \mathcal{L}(\mathbf{L}_{t+1}, \mathbf{c}_t, \Delta \mathbf{p}_t, \mathbf{e}_t, \mathbf{a}_t, \mathbf{B}_t) \quad (6)$$

$$\Delta \mathbf{p}_{t+1} = \underset{\Delta \mathbf{p}_t}{\operatorname{argmin}} \mathcal{L}(\mathbf{L}_{t+1}, \mathbf{c}_{t+1}, \Delta \mathbf{p}_t, \mathbf{e}_t, \mathbf{a}_t, \mathbf{B}_t) \quad (7)$$

$$\mathbf{e}_{t+1} = \underset{\mathbf{e}_t}{\operatorname{argmin}} \mathcal{L}(\mathbf{L}_{t+1}, \mathbf{c}_{t+1}, \Delta \mathbf{p}_{t+1}, \mathbf{e}_t, \mathbf{a}_t, \mathbf{B}_t) \quad (8)$$

Solving for \mathbf{L} : Fixing the $\mathbf{c}_t, \Delta \mathbf{p}_t, \mathbf{e}_t, \mathbf{a}_t$, and \mathbf{B}_t , sub-problem (5) is reduced to:

$$\underset{\mathbf{L}_t}{\operatorname{argmin}} \|\mathbf{L}\|_* + \operatorname{tr}(\mathbf{B}^T (h_2(\mathbf{L}, \mathbf{c}))) + \frac{\mu}{2} \|h_2(\mathbf{L}, \mathbf{c})\|_F^2. \quad (9)$$

The nuclear norm regularized least squared problem (9) has the following closed-form solution:

$$\mathbf{L}_{t+1} = \mathcal{D}_{\frac{1}{\mu_t}} \left[\sum_{i=1}^k \mathcal{R}_{m \times n}(\mathbf{u}_i) c_{i_t} - \mathbf{B}_t / \mu_t \right]. \quad (10)$$

The singular value thresholding (SVT) operator is defined for any matrix \mathbf{Q} with $\mathbf{Q} = \mathbf{U} \Sigma \mathbf{V}^T$ as $\mathcal{D}_\tau[\mathbf{Q}] = \mathbf{U} S_\tau \mathbf{V}^T$ [9], with $S_\tau[\sigma] = \operatorname{sgn}(\sigma) \max(|\sigma| - \tau, 0)$ being the (element-wise) shrinkage operator [10].

Algorithm 1: Solving (4) by the ADM method

Data: Test image \mathbf{X} , initial deformation parameters \mathbf{p} , clean face subspace \mathbf{U} , and the parameter λ
Result: The low-rank clean image \mathbf{L} , the sparse error \mathbf{e} , the coefficient vector \mathbf{c} , and the increments of the deformation parameters $\Delta \mathbf{p}$.

while not converged do

Warp and normalize the image;
 Compute the Jacobian matrix;
 Initialize: $\mathbf{L}_0 = \mathbf{0}, \mathbf{e}_0 = \mathbf{0}, \mathbf{c}_0 = \mathbf{0}, \mathbf{a}_0 = \mathbf{0}, \mathbf{B}_0 = \mathbf{0}$,
 $\mu_0 > 0, \rho > 1$;

while not converged do

Update \mathbf{L}_{t+1} by (10);
 Update $\mathbf{c}_{i,t+1}$ by (12);
 Update $\Delta \mathbf{p}_{t+1}$ by (15);
 Update \mathbf{e}_{t+1} by (17);
 Update the Lagrange multipliers by
 $\mathbf{a}_{t+1} = \mathbf{a}_t + \mu_t (h_1(\Delta \mathbf{p}_{t+1}, \mathbf{c}_{t+1}, \mathbf{e}_{t+1}))$;
 $\mathbf{B}_{t+1} = \mathbf{B}_t + \mu_t (h_2(\mathbf{L}_{t+1}, \mathbf{c}_{t+1}))$;
 Update μ_{t+1} by $\mu_{t+1} \leftarrow \min(\rho \cdot \mu_t, 10^{10})$;
 Check convergence conditions (18) and (19);
 $t \leftarrow t + 1$;

end

$\mathbf{p} \leftarrow \mathbf{p} + \Delta \mathbf{p}$;

end

Solving for \mathbf{c} : Fixing the other variables, sub-problem (6) is reduced to:

$$\begin{aligned} & \underset{\mathbf{c}_t}{\operatorname{argmin}} \mathbf{a}_t^T (h_1(\Delta \mathbf{p}_t, \mathbf{c}_t, \mathbf{e}_t)) + \operatorname{tr}(\mathbf{B}^T (h_2(\mathbf{L}_{t+1}, \mathbf{c}_t))) \\ & + \frac{\mu_t}{2} (\|h_1(\Delta \mathbf{p}_t, \mathbf{c}_t, \mathbf{e}_t)\|_2^2 + \|h_2(\mathbf{L}_{t+1}, \mathbf{c}_t)\|_F^2). \end{aligned} \quad (11)$$

(11) is a quadratic problem which for each $c_i, i \in \{1, \dots, k\}$ admits a closed form solution given by:

$$\mathbf{c}_{i,t+1} = \frac{\mathbf{a}_t^T \mathbf{u}_i + \text{tr}(\mathbf{B}_t^T R_{m \times n}(\mathbf{u}_i))}{2\mu_t} + \frac{\hat{\mathbf{x}}^T \mathbf{u}_i + \text{tr}(\mathbf{L}_{t+1}^T R_{m \times n}(\mathbf{u}_i))}{2}, \quad (12)$$

where $\hat{\mathbf{x}} = \mathbf{x}(\mathbf{p}) + \mathbf{J}(\mathbf{p})\Delta\mathbf{p}_t - \mathbf{e}_t$.

Solving for $\Delta\mathbf{p}$: Sub-problem (7) is written as:

$$\underset{\Delta\mathbf{p}_t}{\text{argmin}} \quad \mathbf{a}_t^T (h_1(\Delta\mathbf{p}_t, \mathbf{c}_{t+1}, \mathbf{e}_t)) + \frac{\mu}{2} \|h_1(\Delta\mathbf{p}_t, \mathbf{c}_{t+1}, \mathbf{e}_t)\|_2^2. \quad (13)$$

By exploiting the fact that \mathbf{U} is orthonormal, each part of (13) is decomposed into the term projected in the $\mathbf{U}\mathbf{U}^T$ and the term projected into orthogonal complement $\mathbf{I} - \mathbf{U}\mathbf{U}^T$. The update of the $\Delta\mathbf{p}$ is obtained by minimizing the projected into orthogonal complement part, i.e.,

$$\underset{\Delta\mathbf{p}_t}{\text{argmin}} \quad \mathbf{a}_t^T (\mathbf{I} - \mathbf{U}\mathbf{U}^T)(\mathbf{x}(\mathbf{p}_t) + \mathbf{J}(\mathbf{p})\Delta\mathbf{p}_t - \mathbf{e}_{t+1}) + \frac{\mu}{2} \|\mathbf{x}(\mathbf{p}) + \mathbf{J}(\mathbf{p}_t)\Delta\mathbf{p}_t - \mathbf{e}_{t+1}\|_{2, \mathbf{I} - \mathbf{U}\mathbf{U}^T}^2. \quad (14)$$

The solution of (14) is given by:

$$\Delta\mathbf{p}_{t+1} = -(\tilde{\mathbf{J}}(\mathbf{p})^T \tilde{\mathbf{J}}(\mathbf{p}))^{-1} \tilde{\mathbf{J}}(\mathbf{p})^T (\mathbf{x}(\mathbf{p}) - \mathbf{e}_{t+1}), \quad (15)$$

where $\tilde{\mathbf{J}}(\mathbf{p})$ is the projected Jacobian in $\mathbf{I} - \mathbf{U}\mathbf{U}^T$. To calculate efficiently the term $\tilde{\mathbf{J}}^T \tilde{\mathbf{J}}$ the following formulation is used: $\tilde{\mathbf{J}}^T \tilde{\mathbf{J}} = \mathbf{J}^T (\mathbf{I} - \mathbf{U}\mathbf{U})^T \mathbf{J}^T = (\mathbf{U}^T \mathbf{J})^T (\mathbf{U}^T \mathbf{J})$.

Solving for \mathbf{e} : Using $\mathbf{L}_{t+1}, \mathbf{c}_{t+1}, \Delta\mathbf{p}_{t+1}$ (8) is written as:

$$\underset{\mathbf{e}_t}{\text{argmin}} \quad \lambda \|\mathbf{e}\|_1 + \mathbf{a}_t^T (h_1(\Delta\mathbf{p}_{t+1}, \mathbf{c}_{t+1}, \mathbf{e}_t)) + \frac{\mu}{2} \|h_1(\Delta\mathbf{p}_{t+1}, \mathbf{c}_{t+1}, \mathbf{e}_t)\|_2^2. \quad (16)$$

The closed-form solution of (16) is given by applying element-wise the shrinkage operator onto: $\mathbf{x}(\mathbf{p}_t) + \mathbf{J}(\mathbf{p})\Delta\mathbf{p}_{t+1} - \mathbf{U}\mathbf{c}_{t+1} + \mathbf{a}_t/\mu_t$, namely

$$\mathbf{e}_{t+1} = \mathcal{S}_{\frac{\lambda}{\mu_t}} [\mathbf{x}(\mathbf{p}_t) + \mathbf{J}(\mathbf{p})\Delta\mathbf{p}_t - \mathbf{U}\mathbf{c}_t + \mathbf{a}_t/\mu_t]. \quad (17)$$

Convergence criteria: The inner loop of the Algorithm 1 terminates when

$$\max(\|\mathbf{e}_t - \mathbf{e}_{t-1}\|_2 / \|\mathbf{x}(\mathbf{p})\|_2, \|\mathbf{L}_t - \mathbf{L}_{t-1}\|_F / \|\mathbf{x}(\mathbf{p})\|_2) \leq \epsilon_2, \quad (18)$$

and

$$\max(\|h_1(\Delta\mathbf{p}_{t+1}, \mathbf{c}_{t+1}, \mathbf{e}_{t+1})\|_2 / \|\mathbf{x}(\mathbf{p})\|_2, \|h_2(\mathbf{L}_{t+1}, \mathbf{c}_{t+1})\|_F / \|\mathbf{x}(\mathbf{p})\|_2) \leq \epsilon_3. \quad (19)$$

The Algorithm 1 terminates when the change of the $\|\mathbf{L}\|_* + \lambda \|\mathbf{E}\|_1$ between two successive iterations is smaller than the threshold ϵ_1 or the maximum number of the outers' loop iterations is reached. The dominant cost of each iteration of Algorithm 1 is that of the SVD algorithm involved in the computation of the SVT operator in update of \mathbf{L} . Consequently, the computational complexity of Algorithm 1 is $\mathcal{O}(T(\min(m, n)^3 + n^2 k))$, where T is the total number of iterations until convergence.

In Fig. 4, the convergence of the inner loop of Algorithm 1 is depicted. The low-rank and error images produced after 30, 50 and 117 iterations, respectively, are also shown.




Fig. 4. The convergence curve of the Algorithm's inner loop.

IV. EXPERIMENTAL RESULTS

The performance of the FAR is assessed in: a) frontal face reconstruction, b) landmark localization, and c) pose invariant face recognition and verification, by conducting experiments in 6 facial image databases, which are described briefly next.

A. Data description

LFPW: The Labeled Faces Parts in-the-wild (LFPW) [6] database contains images downloaded from the internet (i.e., google.com, flickr.com etc), images exhibiting multiple variations such as pose, expression, illumination, and occlusions. Since only the URLs of images were provided, 811 out of the 1,132 training images and 224 out of the 300 test images were downloaded.

HELEN: The HELEN [22] database consists of 2,300 images downloaded from Flickr web service, containing a broad range of appearance variation, including pose, lighting, expression, occlusion, and individual differences. The size of the face in each of the images was approximately 500×500 pixels.

AFW: The Annotated Faces in-the-wild (AFW) [44] database consists of 250 images with 468 faces. That is more than one faces are annotated in each image. The images exhibit similar variations with those in the LFPW and HELEN databases.

FERET: The Facial Recognition Technology (FERET) [28] database consists of 14,051 images of 200 different subjects. All images capture the same 'Neutral' expression for 9 different poses under different illuminations, where each subject also has an additional image with a random facial expression.

MultiPIE: The CMU Multi Pose Illumination and Expression (MultiPIE) [17] database consists of approximately 750,000 images from 337 subjects, captured under 6 different expressions, 15 poses, and 19 illuminations.

LFW: The Labeled Faces in the Wild (LFW) [20] database contains 13,233 images of 5,749 people downloaded from

the Web and is designed as a benchmark for the problem of unconstrained automatic face verification. All images are characterized by the existence of large pose, expression and occlusion variations.

B. Experimental setup

In all experiments, the orthonormal clean face subspace \mathbf{U} was constructed by employing only frontal view without occlusions face images. In total 500 frontal images (217 from the training set of the LFPW and 283 from the training set of the HELEN databases) were selected to build the bases \mathbf{U} . The frontal images were warped in a common frame (185×193 pixels) by using a piece-wise affine motion model and subsequently the PCA was applied on the warped shape-free textures. The first $k = 450$ eigen-images were kept. Unless otherwise stated, throughout the experiments, the same \mathbf{U} was used and the parameters of the Algorithm 1 were fixed: $\lambda = 0.3$, $\rho = 1.1$, $\mu_0 = 10^{-6}$, $\epsilon_1 = 10^{-3}$, $\epsilon_2 = 10^{-5}$, and $\epsilon_3 = 10^{-7}$.

C. Frontal face reconstruction

Next, the ability of the FAR to reconstruct frontal faces from non-frontal images of unseen subjects is investigated by using two unseen subjects taken from MultiPIE and from FERET databases and 5 in-the-wild images.

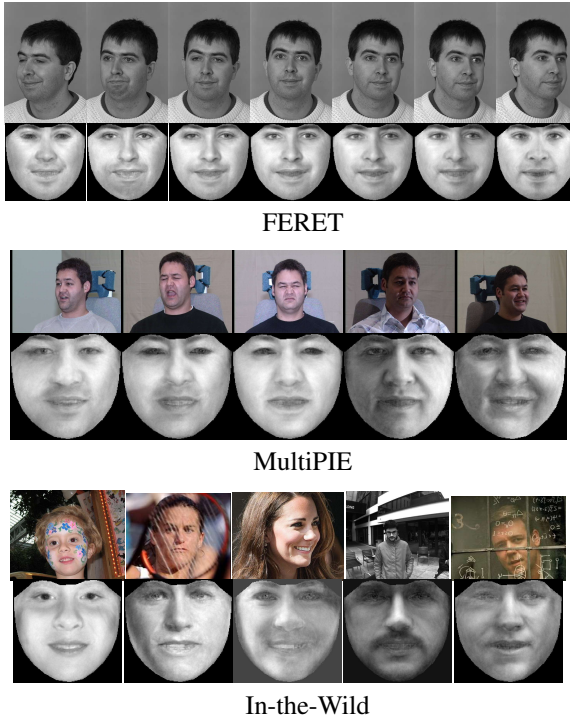


Fig. 5. Reconstructed frontal images of unseen subjects under controlled and in-the-wild conditions.

Given an input facial image, the initialization was produced by applying the detector [44]. The image, the initialization and \mathbf{U} were given as input into Algorithm 1. By using the produced $\Delta \mathbf{p}$ the outer loop was rerun again for one iteration without updating $\Delta \mathbf{p}$ the execution of sub-problem (15).

In Fig. 5 (rows: 1-2) the reconstructed frontal faces from the non-frontal images ('ba', 'bc', 'bd', 'be', 'bf', 'bg', and 'bh') of '00268' subject from FERET database are illustrated. Fig. 5 (rows: 3-4) depicts the frontal reconstructed views from the images taken from MultiPIE with (a) 'Surprise' at -30° , (b) 'Scream' at -15° , (c) 'Squint' at 0° , (d) 'Neutral' at $+15^\circ$, and (e) 'Smile' at $+30^\circ$. By visually inspecting Fig. 5, it is clear that the FAR is robust to pose, expression, and lighting conditions variations. This attributed to the fact that the matrix ℓ_1 -norm was adopted for non-Gaussian noise characterization. Frontal reconstructed views from in-the-wild images are depicted in Fig. 5 (rows: 5-6).

To quantitatively assess the quality of the frontalized images the following experiment was conducted. To this end, 'Neutral' images of 20 different subjects of MultiPIE under poses $-30^\circ : 30^\circ$ (5 for each subject, 100 in total) were selected. The images of each subject were frontalized by employing the FAR. The Root Mean Square Error (RMSE) between each frontalized image and the real frontal image of the subject is used as evaluation metric. The performance of the FAR with respect to RMSE is compared with that obtained by the frontalization system of the DeepFace [34]. The average RMSEs of the FAR and DeepFace are 0.0817 and 0.1025, respectively. It is worth noting that, even DeepFace employs a 3D model to handle out-of-planar rotations, the FAR performs better without using any kind of 3D information.

D. Face landmark localization

The performance of the FAR in the generic face alignment problem is assessed by conducting experiments on in-the-wild databases namely. the LFPW, the HELEN and the AFW. To this end, the performance of the FAR is compared against that obtained by (a) the AAMs, the CLMs, and the SDM using exactly the same training data as well as the same features and (b) the state-of-the-art method and features. The annotations provided in [30], [31] have been employed for evaluation purposes. The average point-to-point Euclidean distance of 49 interior landmark points (excluding the points correspond to face boundary) normalized by the Euclidean distance of the outer corner of eyes is used as the evaluation measure. In addition, the cumulative error distribution curve (CED) for each method is computed by using the fraction of test images for which the average error was smaller than a threshold.

Same train set and features: In order to compare fairly the competing methods, the same training data, initialization, and features were employed. The 500 frontal images used to build the \mathbf{U} , were used as the training set while the pixel intensities (PIs) were selected as the texture representation. The results produced by the detector [44] were used to initialize the methods. For this experiment the implementations provided by the platform MENPO [1] were used for all methods.

The CEDs produced by all methods for the LFPW (test set), the HELEN (test set), and the AFW databases are depicted in Fig. 6. Clearly, the FAR outperforms the AAMs-PIs, the CLMs-PIs, and the SDM-PIs. More specifically, for normalized error equal to 0.05³ the FAR yield an 20.1%, 21.5% and 24.6% improvement compared to that obtained by the AAMs-PIs

³This value was found by visually inspecting the results.




Fig. 6. Cumulative error distribution curves produced by the CLMs-PIs, the AAMs-PIs, the SDM-PIs, and the FAR for the LFPW, the HELEN, and the AFW.




Fig. 7. Comparison of the cumulative error distribution curves obtained by the FAR and the SDM on the LFPW, the HELEN, and the AFW databases.



AAMs-PIs CLMs-PIs SDM-PIs SDM-SIFT FAR-PIs

Fig. 8. Sample fitting results from the LFPW (rows: 1-2), the HELEN (rows: 3-4), and the AFW (rows: 5-6) databases. (The Figure is best viewed in color)

across the test databases, respectively. A few fitting examples from the test databases are depicted in Fig. 8.

State-of-the-art method and features:

In this experiment, the FAR is compared against the state-

of-the-art method SDM [39]. In particular, the implementation as well as the pre-trained model provided by the authors were used. Both the FAR and the SDM were initialized by using the same detector [44], and the SDM was initialized following the instructions of the authors⁴. The CEDs from this experiment are shown in Fig.7 where the FAR achieves comparable performance with that obtained by SDM using only **500 frontal images**. It is worth mentioning, the SDM was trained on thousand images captured under several variations including different poses, illuminations and expression. Furthermore, the SDM method takes full advantage of SIFT - a powerful hand-crafted feature - while the FAR employs only pixel intensities.

E. Pose-invariant face recognition and verification

The performance of the FAR on pose invariant face recognition and verification is assessed by conducting experiments on the MultiPIE, FERET, and LFW databases. Apart from the experimental results reported in this section more can be found on the supplementary materials.

1) *Pose invariant face recognition*: The frontal views of all images were reconstructed following the methodology described in Section IV-C. The reconstructed images were cropped to remove the surrounding black pixels. The Image Gradient Orientations (IGOS) features [36] were used for image representation. The dimensionality of IGOs was reduced by applying PCA. The classifier in [41] was used.

The performance of the FAR is compared against of that obtained by the following 2D based methods LGBP [42] and PIMRF [18], 3D based methods 3DPN [4], EGFC [24], and PAF [40] as well as the Deep learning based methods SPAE [21] and DIPFS [45]. It should be noticed that all methods were evaluated under the fully automatic scenario; where both

⁴<http://www.humansensing.cs.cmu.edu/intraface>

the bounding box of the face region and the facial landmark were located automatically.

Results on FERET: One frontal image, denoted as ‘ba’, from each of the 200 subjects was used to form the gallery set, while the images captured at 6 different poses i.e., from -40° to 40° were selected as the probe images.

TABLE I. RECOGNITION RATES (%) ACHIEVED BY THE COMPARED METHODS ON THE FERET DATABASE.

Method	bh -40°	bg -25°	bf -15°	be $+15^\circ$	bd $+25^\circ$	bc $+40^\circ$	Avg
LGBP [42]	90.5%	98.0%	98.5%	97.5%	97.0%	91.9%	95.6%
3DPN [4]	90.5%	98.0%	98.5%	97.5%	97.0%	91.9%	95.6%
PIMRF [18]	91.0%	97.3%	98.0%	98.5%	96.5%	91.5%	95.5%
PAF [40]	98.0%	98.5%	99.25%	99.25%	98.5%	98.0%	98.56%
FAR	96.5%	99.0%	100.0%	100.0%	100%	96%	98.58%

In Table I the recognition rates achieved by the competing methods in the different poses are reported. Clearly, the FAR (recognition accuracy 98.58%) outperforms both the 2D and 3D state-of-the-art methods PIMRF and PAF, respectively. It is worth mention that the PIMRF employs 200 images from the FERET database (different from the test set) in order to train the frontal synthesizer. Consequently, the different lighting conditions of the database are taken into account. This is not the case for the FAR where only frontal images from a generic in-the-wild database (i.e., the LFW and HELEN) have been used. Even the FAR does not use any kind of 3D information it outperforms the PAF where an elaborated 3D model (trained from 4.624 facial scans) has been used to find the 3D pose and extract pose adaptive features. The reported results of the EGFC [24] were not included in Table I as they were obtained using a semi-automatic protocol (i.e., 5 manually annotated landmark points used).

Results on MultiPIE: The images of the 137 subjects (Subject ID 201: 346) with ‘Neutral expression and poses $-30^\circ : +30^\circ$ captured under 4 different sessions were selected. The gallery was created by the frontal images of the earliest session for each subject, while the rest images including frontal and non-frontal views were used as probes. It should be mentioned that images of first 200 subjects which include all poses (4207 in total) were used for training purposes. In particular, the above mentioned images were used in the 3DPN to train view-based models, in the SPAE, DIPFS to train the deep networks, and in the EGFC to train the pose estimator and matching model. The recognition accuracy achieved by the just mentioned methods is reported in Table II. Again, the FAR outperforms four out of five methods that is compared to, verifying the high quality of the frontalized images. The FAR also performs comparable with the DIPFS by simply using just 500 frontal images outside the MultiPIE.

2) *Face verification:* The performance of the FAR in the face verification under in-the-wild conditions is assessed by conducting experiment in the LFW database, using the ‘image-restricted, no outside data results’ protocol. The reported results are obtained using 10-fold cross validation.

In this experiment the basis U and the detector [44] were not used since they based on images outside the database. To create the initializations and a new U_{LFW} , the method for automatic construction of deformable models presented in [2]

TABLE II. RECOGNITION RATES (%) ACHIEVED BY THE COMPARED METHODS ON THE MULTIPLE DATABASE.

Method	130_06 -30°	140_06 -15°	051_07 0°	050_08 15°	041_08 30°	Avg
PIMRF [18]	89.7%	91.7%	92.5%	91.0%	89.0%	90.78%
3DPN [4]	91.0%	95.7%	96.9%	95.7%	89.5%	93.76%
SPAE [21]	92.6%	96.3%	-	95.7%	94.3%	94.72%
EGFC [24]	95.0%	99.3%	-	99.0%	92.9%	96.55%
DIPFS [45]	98.5%	100%	-	99.3%	98.5%	99.07%
FAR	94.3%	98.7%	99.4%	97.3%	95.6%	97.06%

was employed. The goal of this method is to build a deformable model using only a set of images with the corresponding bounding boxes. To define the bounding boxes without using a pre-trained detector, the deep funneled images of the LFW [19] were employed. Therefore, since these images are aligned the exact bounding box is known. Subsequently, a deformable model was built automatically from the training images of each fold. The created model was fitted to all images and those (from training images) with fitted shapes similar to mean shape were selected to build the bases U_{LFW} . In each fold the images were frontalized using the U_{LFW} and they were cropped next. In the sequel, the gradient orientations ϕ_1, ϕ_2 of each image pair were extracted and the cosine of difference between them $\Delta\phi = \phi_1 - \phi_2$ normalized to the range $[0 - 2\pi]$, was used as the feature of the pair. These features are classified by a Support Vector Machine (SVM) with an RBF kernel. The performance of the FAR is compared against that obtained by the MRF-MLBP [3], Fisher Vector Faces [33] and the EigenPEP [23] methods⁵. The mean classification accuracy and the corresponding standard deviation on LFW are reported in Table III. By inspecting Table III the FAR outperforms the MRF-MLBP and the Fisher Vector Faces and performance comparably with the recently published method EigenPEP.

MRF-MLBP [3]	0.7908 ± 0.0014
Fisher vector faces [33]	0.8747 ± 0.0149
EigenPEP [23]	0.8897 ± 0.0132
FAR	0.8881 ± 0.0078

TABLE III. MEAN CLASSIFICATION ERROR AND STANDARD DEVIATION ON THE LFW DATABASE.

V. CONCLUSIONS

In this paper we developed the first, to the best our knowledge, method that jointly performs landmark localization and face frontalization using only a simple statistical model of few hundred frontal images. The proposed method outperforms state-of-the-art methods for face landmark localization that were trained on thousands of images in many poses and achieves comparable results in pose invariant face recognition and verification without using 3D elaborate models or Deep Learning-based features extraction.

REFERENCES

- [1] J. Alabort-i Medina, E. Antonakos, J. Booth, P. Snape, and S. Zafeiriou. Menpo: A comprehensive platform for parametric image alignment and visual deformable models. In *Proceedings of the ACM International Conference on Multimedia*, pages 679–682. ACM, 2014. 5

⁵These methods are the 3 top performing in the LFW according to <http://vis-www.cs.umass.edu/lfw/results>.

- [2] E. Antonakos and S. Zafeiriou. Automatic construction of deformable models in-the-wild. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1813–1820. IEEE, 2014. 7
- [3] S. R. Arashloo and J. Kittler. Efficient processing of mrfs for unconstrained-pose face recognition. In *Biometrics: Theory, Applications and Systems (BTAS), 2013 IEEE Sixth International Conference on*, pages 1–8. IEEE, 2013. 7
- [4] A. Athana, T. K. Marks, M. J. Jones, K. H. Tieu, and M. Rohith. Fully automatic pose-invariant face recognition via 3d pose normalization. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 937–944. IEEE, 2011. 1, 6, 7
- [5] A. Athana, S. Zafeiriou, S. Cheng, and M. Pantic. Robust discriminative response map fitting with constrained local models. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3444–3451. IEEE, 2013. 1
- [6] P. N. Belhumeur, D. W. Jacobs, D. Kriegman, and N. Kumar. Localizing parts of faces using a consensus of exemplars. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 545–552. IEEE, 2011. 4
- [7] D. P. Bertsekas. Constrained optimization and lagrange multiplier methods. *Computer Science and Applied Mathematics, Boston: Academic Press*, 1982, 1, 1982. 3
- [8] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(9):1063–1074, 2003. 1
- [9] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010. 3
- [10] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11, 2011. 3
- [11] X. Cheng, C. Fookes, S. Sridharan, J. Saragih, and S. Lucey. Deformable face ensemble alignment with robust grouped-l1 anchors. In *Automatic Face and Gesture Recognition (FG), 2013 10th IEEE International Conference and Workshops on*, pages 1–7. IEEE, 2013. 2
- [12] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. *IEEE Transactions on pattern analysis and machine intelligence*, 23(6):681–685, 2001. 1
- [13] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995. 1
- [14] D. Cristinacce and T. F. Cootes. Feature detection and tracking with constrained local models. In *BMVC*, volume 2, page 6, 2006. 1
- [15] D. Donoho. For most large under-determined systems of linear equations, the minimal l1-norm near-solution approximates the sparsest near-solution. *Communications on Pure and Applied Mathematics*, 59(7):907–934, 2006. 3
- [16] M. Fazel. *Matrix rank minimization with applications*. PhD thesis, PhD thesis, Stanford University, 2002. 3
- [17] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010. 4
- [18] H. T. Ho and R. Chellappa. Pose-invariant face recognition using markov random fields. *Image Processing, IEEE Transactions on*, 22(4):1573–1584, 2013. 1, 6, 7
- [19] G. Huang, M. Mattar, H. Lee, and E. G. Learned-Miller. Learning to align from scratch. In *Advances in Neural Information Processing Systems*, pages 764–772, 2012. 7
- [20] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 2, 4
- [21] M. Kan, S. Shan, H. Chang, and X. Chen. Stacked progressive auto-encoders (spae) for face recognition across poses. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1883–1890. IEEE, 2014. 6, 7
- [22] V. Le, J. Brandt, Z. Lin, L. Bourdev, and T. S. Huang. Interactive facial feature localization. In *Computer Vision-ECCV 2012*, pages 679–692. Springer, 2012. 4
- [23] H. Li, G. Hua, X. Shen, Z. Lin, and J. Brandt. Eigen-pep for video face recognition. *Submitted to ECCV and under review*, 2104, 2014. 7
- [24] S. Li, X. Liu, X. Chai, H. Zhang, S. Lao, and S. Shan. Morphable displacement field based image matching for face recognition across pose. In *Computer Vision-ECCV 2012*, pages 102–115. Springer, 2012. 6, 7
- [25] I. Matthews and S. Baker. Active appearance models revisited. *International Journal of Computer Vision*, 60(2):135–164, 2004. 1
- [26] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995. 3
- [27] Y. Peng, A. Ganesh, J. Wright, W. Xu, and Y. Ma. Rasl: Robust alignment by sparse and low-rank decomposition for linearly correlated images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(11):2233–2246, 2012. 2
- [28] P. J. Phillips, H. Moon, S. A. Rizvi, and P. J. Rauss. The feret evaluation methodology for face-recognition algorithms. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(10):1090–1104, 2000. 4
- [29] C. Sagonas, Y. Panagakis, S. Zafeiriou, and M. Pantic. Raps: Robust and efficient automatic construction of person-specific deformable models. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1789–1796. IEEE, 2014. 2
- [30] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *Computer Vision Workshops (ICCVW), 2013 IEEE International Conference on*, pages 397–403. IEEE, 2013. 1, 5
- [31] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. A semi-automatic methodology for facial landmark annotation. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2013 IEEE Conference on*, pages 896–903. IEEE, 2013. 1, 5
- [32] J. M. Saragih, S. Lucey, and J. F. Cohn. Deformable model fitting by regularized landmark mean-shift. *International Journal of Computer Vision*, 91(2):200–215, 2011. 1
- [33] K. Simonyan, O. M. Parkhi, A. Vedaldi, and A. Zisserman. Fisher vector faces in the wild. In *Proc. BMVC*, volume 1, page 7, 2013. 2, 7
- [34] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1701–1708. IEEE, 2014. 1, 5
- [35] G. Tzimiropoulos, J. Alabort-i Medina, S. Zafeiriou, and M. Pantic. Generic active appearance models revisited. In *Computer Vision-ACCV 2012*, pages 650–663. Springer, 2013. 1
- [36] G. Tzimiropoulos, S. Zafeiriou, and M. Pantic. Subspace learning from image gradient orientations. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(12):2454–2466, 2012. 6
- [37] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996. 3
- [38] P. Viola and M. Jones. Rapid object detection using a boosted cascade of simple features. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages 1–511. IEEE, 2001. 1
- [39] X. Xiong and F. De la Torre. Supervised descent method and its applications to face alignment. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 532–539. IEEE, 2013. 1, 6
- [40] D. Yi, Z. Lei, and S. Z. Li. Towards pose robust face recognition. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3539–3545. IEEE, 2013. 1, 6, 7
- [41] D. Zhang, M. Yang, and X. Feng. Sparse representation or collaborative representation: Which helps face recognition? In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 471–478. IEEE, 2011. 6
- [42] W. Zhang, S. Shan, W. Gao, X. Chen, and H. Zhang. Local gabor binary pattern histogram sequence (lgbphs): A novel non-statistical model for face representation and recognition. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 1, pages 786–791. IEEE, 2005. 6, 7
- [43] Z. Zhang, A. Ganesh, X. Liang, and Y. Ma. Tilt: transform invariant low-rank textures. *International Journal of Computer Vision*, 99(1):1–24, 2012. 2
- [44] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2879–2886. IEEE, 2012. 4, 5, 6, 7
- [45] Z. Zhu, P. Luo, X. Wang, and X. Tang. Deep learning identity-preserving face space. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 113–120. IEEE, 2013. 7